Tracy Slatyer – Lecture 1 Questions

Questions marked in green were answered during the Q&A session. I haven't tried to correct grammar/spelling. Where a slide number was given it is shown.

Q1 (slide 5). Why does DM not need to be weak-scale if it annihilates?

A: Annihilation is a ubiquitous process for DM candidates that are their own antiparticles (Majorana fermions, real scalars, etc) - if no symmetry forbids it, it will generically occur. Of course that's a different question than whether annihilation sets the relic abundance. For DM that freezes out via two-body annihilation while non-relativistic, the abundance ends up being roughly independent of the DM mass and instead is fixed by the annihilation cross section. So any model with correct annihilation cross section will get roughly the right relic abundance. One way to get that correct annihilation cross section is to have weak-scale DM with weak-scale interactions annihilating at tree level (so the rate scales roughly as coupling^2/mass^2), but you can also have much lighter DM with much weaker couplings, as a simple example. In more complicated examples, there's no reason the rate has to be given just by coupling^2/mass^2 - you can have a phase space suppression, for example, as in Forbidden Dark Matter https://arxiv.org/abs/1505.07107 .

Q2 (slide 5). How much does the "coldness" or "warmness" of DM particles today effect the thermal relic benchmark?

A: There are a few things this question could mean - I tried to answer a few of them in real-time, but let me try to be a bit more complete here.

-Most of the time when we think about thermal relic dark matter we're thinking about DM that is MeV-scale and heavier, as lighter thermal relic dark matter is typically excluded e.g. by bounds from Big Bang Nucleosynthesis on the number of relativistic degrees of freedom. In contrast, when we think about dark matter that is fast-moving enough to be "warm" in an observational sense, we are usually thinking about keV-scale dark matter (the current bounds on warm dark matter probe masses around 5 keV if I remember correctly). So usually "warm dark matter" and "thermal relic dark matter" are somewhat mutually exclusive - we usually think about alternative production mechanisms for warm dark matter.

-Dark matter that freezes out while relativistic has very different behavior to the cold dark matter thermal freezeout scenario - there's no exponential suppression to the dark matter density, and this in turn means that the dark matter needs to be down around the 1 eV scale in mass to avoid overclosing the universe (independent of the annihilation cross section), which is ruled out for multiple reasons. Warm dark matter scenarios are in between the hot and cold dark matter extremes, but if the dark matter is much heavier than 1 eV then it would need its abundance to be exponentially depleted if we assume the standard cosmology, i.e. the non-relativistic approximation would probably not be terrible right at freezeout. Assuming non-relativistic freezeout and fixing the level of exponential depletion to match the observed abundance today gives you an estimate for the relic density that is only logarithmically dependent on the mass. So I would expect the thermal relic cross section for keV-scale dark matter (which is “warm” in that it doesn’t need nearly as much exponential suppression of its abundance to match the observed relic density, compared to GeV dark matter) undergoing standard thermal freezeout to be somewhat different from GeV-scale dark matter, but not by more than an order of magnitude or two. I don’t know off-hand of a reference for the exact thermal cross section at these low masses - here is a careful thermal relic calculation: https://arxiv.org/pdf/1204.3622.pdf, but it only goes down to 100 MeV - but again it’s generically quite challenging to evade constraints from Big Bang Nucleosynthesis on thermal dark matter lighter than a MeV (although here is one idea for doing so https://arxiv.org/abs/1706.07046).

-The “coldness” or “warmness” of dark matter today could change the structure of dark matter halos on small scales, which in turn could affect the size of annihilation signals. In this case the modification would be to the J-factor, not the cross section.

-In models where <sigma v> is velocity-dependent, the “warmness” or “coldness” of the dark matter could directly affect <sigma v>, but (a) the default assumption is usually velocity-independent <sigma v>, (b) once dark matter has collapsed into haloes, its velocity is typically primarily set by the halo gravitational potential, not any primordial “warmness”.

Q3 (slide 5). How thermal average cross-section inversely depends on the matter-radiation equality temperature, I thought Gamma = H will give, <sigmaV> inversely proportional to freeze-out temperature times Planck mass. Can you please elaborate?

A: Sure! The argument runs as follows:
Freezeout occurs when n<sigma v> \sim H
If freezeout occurs during radiation domination, we can write H_f \sim T_f^2/m_Planck, where “f” subscripts indicate “at freezeout”, so we have n_f <sigma v> \sim T_f^2/m_Planck
We want n_f to correspond to the correct late-time relic density, which is roughly equivalent to requiring that we get the temperature of matter-radiation equality right (since matter-radiation equality describes the amount of energy density in matter, which is mostly dark matter).
Thus (using “MRE” subscripts to indicate matter-radiation equality) we want:
(rho_DM)_MRE = m_DM n_MRE \approx (\rho_rad)_MRE \sim T_MRE^4 (here I’m ignoring all O(1) numerical factors, and equating the dark matter and radiation densities at MRE)
Now if freezeout occurred during radiation domination, n_MRE and n_f are just related by the expansion of the universe between freezeout and MRE, and so (ignoring changes in the number of effective degrees of freedom) we can roughly write n_MRE \approx (T_MRE/T_f)^3 n_f .
Then we have m_DM (T_MRE/T_f)^3 n_f \approx T_MRE^4, so to get the correct relic density requires achieving a freezeout number density of:
n_f \approx T_MRE T_f^3 / m_DM

But then from n_f <sigma v> \sim T_f^2/m_Planck , we can eliminate n_f and find that the cross section needed to get the right relic density satisfies:

<sigma v> \sim T_f^2/m_Planck / (T_MRE T_f^3/m_DM) = (m_DM/T_f) / (T_MRE m_Planck)

So you’re right there’s a factor of 1/T_f, but T_f scales with m_DM (because cold dark matter only freezes out when its number density starts to get exponentially suppressed due to T dropping below m_DM), and so this dependence mostly cancels out. The cancellation between m_DM and T_f is likewise why there’s only a weak (logarithmic) dependence on the DM mass in the thermal cross section.

Note that this cancellation doesn’t hold for 3-body or higher annihilation processes, where consequently the equivalent of the thermal relic cross section scales strongly with mass - it’s a good exercise to follow the logic through in this case and see what happens. (Basically you have extra powers of n_f in the equation relating n_f to H_f, and this leads to extra powers of m_DM.)

Q4 (slide 7). Could you briefly explain again the order of magnitude calculations you did at the bottom of the slide - bit unclear to me where those expressions are from

A: If I’m looking at the right slide, I think this is degenerate with Q7 below, I’ll put the answer there.

Q5 (slide 8). What if W/Z bosons are produced? Higgs?

A: These have large branching ratios into hadronic decay modes, and consequently the statements for hadronic channels apply (i.e. the strongest limits are usually from photon and antiproton searches).

Q6 (slide 10). In general, how do we tell a redshifted photon from one that just started out at a lower frequency?

A: If you have a feature in the spectrum, such as an emission or absorption line, you can look for the lineshape. If you have good reason to think the photon is coming from a particular object (e.g. a distant galaxy cluster), and you can measure the distance to that object by an independent method, then you can infer the redshift and hence the original energy. But yes, in general you need additional information beyond just “I have measured a photon at a specific energy”.

Q7 (slide 7). Except for the DM decay time having to be inversely proportional to some mass scale dimensionally, what determines the powers of the GUT and DM mass scales in the higher dimension operators?

A: So first, don’t take this estimate too seriously, this is just a quick back-of-the-envelope calculation. There is no guarantee at all that there’s an extra physics scale related to dark matter at the GUT scale - I just wanted to give a bit of a framework for what kinds of physics controlling decays might be excluded, detectable, or invisible.

Imagine we can express the decay as a tree-level process determined by some effective operator in the Lagrangian (which can be dimension 5, 6, 7, etc), consisting only of low-mass fields (which includes the DM in this case). The physics at the high scale (in this case M_GUT) has been integrated out to produce this effective operator. A dimension-5 operator has total mass dimension 5 for the fields, but all Lagrangian terms must have mass dimension 4, so the operator must have a coupling/prefactor with dimension mass^(-1). When I say “a dimension-5 operator suppressed by a high scale M” (e.g. M=m_GUT), I mean that coupling is parametrically 1/M. The tree-level decay rate controlled by this operator will scale like the coupling^2, i.e. like 1/M^2, and so the lifetime will scale as M^2 (i.e. the inverse of the rate). For dimension-6 it will be M^4, for dimension 7 M^6, etc. The lifetime needs to have units mass^(-1) overall, so we will make the dimension correct with powers of the low-scale mass (which will come from the low-mass fields making up the operator). This gives the scalings on slide 7.

(Note also that I am ignoring all numerical factors here.)

Q8 (slide 11). How safe is it to assume that the dark matter annihilates isotropically?

A: At the level of an individual interaction, there’s no reason to think the annihilation products should be emitted isotropically. However if the dark matter is virialized/equilibrated, usually the velocity distribution will be roughly isotropic (e.g. in the Galactic frame), and so when we integrate over a large number of dark matter collisions there should be no preferred direction for the annihilation products. This reasoning might not work if there was a high-velocity non-equilibrium dark matter component that was responsible for most annihilations.

Q9 (slides 17, 18). What are we setting the expression at the end of slide 17 equal to, to get the expression at the start of slide 18?

A: I assume this is referencing the pdf page numbers rather than the slide numbers. If this is the case, the expression at the end of slide 17 is equal to dN_obs/dE dt (i.e. the dN/dE spectrum per unit time), from the volume element dV, with a detector of area A. The expression at the top of slide 18 is obtained by replacing A -> dA to consider an infinitesimal detector element, dividing both sides by dA, and integrating over dV (dV = R^2 dR d\Omega, this is just the volume element in spherical polar coordinates).

Q10 (slide 12). Where did the 1/R^2 term go that we had in the last eq on s11?

A: The volume element in spherical polar coordinates is dV = R^2 dR d\Omega - sorry, I should’ve pointed this out in words, the R^2 factor in dV absorbs the 1/R^2 from slide 11.

Q11 (slide 18). Could you motivate a little bit more where we get the J-factor? I'm afraid I don't see how it comes in / how it connected to the previous slide.

A: Ah, I suspect this may have been a matter of someone quoting pdf page numbers rather than slide numbers. If this was about the slide numbered 12 (page 18 in the pdf), hopefully the two answers above help resolve the question.

If it was about the slide numbered 18, the mention of the J-factor there is in the context of dwarf galaxy J-factors being uncertain, and hence inducing an uncertainty on the cross section we can constrain with dwarfs. There are a couple of recent papers cited on the slide, which explored the uncertainties in both line-of-sight foregrounds and J-factors for dwarf analyses, and found them to be significant. In particular, the J-factor is hard to measure because dwarf galaxies (especially ultrafaint ones) have relatively few stars associated with them, which means we don't have many tracers with which to measure the dark matter distribution. (Lack of stars and baryonic matter = low backgrounds, which is beneficial for dark matter searches, but it does have this unfortunate side-effect.)

Q12 (slide 13). What is r_s in the equation on slide 13? Thanks!

A: This equation is the Navarro-Frenk-White profile, which is a dark matter density profile commonly fitted to numerical dark-matter-only simulations (it remains unclear how well it describes actual galaxies when you get very close to the center). r_s is the scale radius parameter, which varies from galaxy to galaxy; for the Milky Way it's about 20 kpc. The key part of the NFW profile, for the purpose of the discussion on the slide, is that it scales as 1/r for r << r_s, i.e. close to the Galactic Center - this is responsible for the large J-factor at the Galactic Center.

Q13 (slide 13). While on slide 13, you mentioned that an actual detection of DM annihilation could tell us whether the density profile is cuspy or cored. But wouldn't we first have to tease apart the J-factor from the particle physics prefactor? How do we do that?

A: Great question! Under standard assumptions (no significant redshifting/absorption, no velocity dependence in the annihilation cross section, etc) the spatial distribution of the signal is completely controlled by the J-factor; the particle physics prefactor is energy-dependent but can be factored out of the line-of-sight integral. So we would not be able to get the absolute *normalization* of the dark matter density by seeing an annihilation signal - that would be degenerate with the cross-section - but by looking at different lines of sight, we would be able to make some inferences about the distribution of the dark matter. If we assumed spherical symmetry for the dark matter halo (a bit of a fishy assumption but likely better toward the Galactic Center), we would be able to infer the radial profile (with a resolution controlled by the resolution of the instrument).

For example, the Galactic Center excess (which I'll discuss a bit next time) appears to have a brightness profile roughly consistent with a spherically symmetric signal of flux/unit volume \propto r^(-2.5), where r is the distance from the center of the Galaxy. If interpreted as a dark matter annihilation signal, this would suggest that the dark matter density scales as r^(-1.25), since the annihilation rate scales as density squared and we are assuming no position dependence in the particle-physics parameters.

Q14 (slide 13). How will we differentiate photons from dark matter to photons from baryonic matter ?

A: This is a very key question! Generally you perform a likelihood analysis (or similar) where you include a model for your dark matter signal on one hand, and your astrophysical backgrounds on the other. The appropriate astrophysical background depends on the search, it can be a theoretical model or something more data-driven (typically at least some freedom to fit the data is required). The more information you have about properties of your dark matter signal, the easier it is to perform this separation - if you're working with photons/neutrinos you can use spatial information, as you know which line-of-sight the particles arrived along, as well as energy information, whereas for charged cosmic rays you typically only have energy information.

That said it is generally not possible to distinguish at the level of individual photons - the best you can hope for is to say (for example) "the best-fit model has 50% of the photons in this bin coming from dark matter and 50% from astrophysics". The exceptions are searches like those for antideuterons and anti-helium, which have the hope of being background-free - I'll discuss this more tomorrow.

Q15. How do things like the bullet cluster constrain the DM annihilation?

A: Individual clusters do not usually give the strongest constraints on DM annihilation (unless you assume a very large amount of dark matter substructure), and at present I don't know of any strong constraints from Bullet Cluster observations. This paper explored the idea of cross-correlating annihilation/decay signals with gravitational lensing to identify the dark matter distribution:
https://arxiv.org/pdf/1502.03824.pdf

If you meant "do limits on dark matter self-interaction place upper bounds on the annihilation rate", they can in the context of some specific models, but in general there's no reason to think the self-interaction rate and the rate of annihilation to standard model particles are closely linked.

If you mean "does the Bullet Cluster bound DM-SM scattering cross sections, and consequently constrain DM-SM annihilation interactions", then usually direct-detection experiments set stronger limits on scattering than astro/cosmo searches, although there can be notable exceptions at low dark matter masses. The relationship between DM-baryon scattering and DM annihilation is also quite model-dependent.

Q16. Have people discussed indirect detection via gravitational waves at all?

A: I know people have thought about it in the context of:
-Black holes as dark matter (e.g. https://arxiv.org/abs/1603.00464)
-Stochastic gravitational wave signals from phase transitions due to modified early-universe cosmology (e.g. https://arxiv.org/abs/1912.06139)
-Searches for axions bound to black holes (e.g. https://arxiv.org/abs/1411.2263)

If DM is just a collection of almost-inert particles peacefully floating around and occasionally annihilating, it’s hard for it to do much in the way of making gravitational waves, but there are certainly models with interesting signatures to be found.

Q17. Can we combine limits set from multiple sources into one stronger limit, or is this statistically impossible since we're using different principles of detection/particle physics for each indirect detection experiment?

A: Yes, the likelihoods can be combined, but you do need to be a bit careful about the systematics. The joint Fermi/HAWC/VERITAS/MAGIC/HESS paper cited on slide 18 is an example of such an analysis (https://arxiv.org/abs/1909.06310).

Q18 (slide 17). Why do they specifically look for annihilation to bs or taus if all the SM particles are possible channels?

A: They’re just examples - often other quark channels look a lot like b’s, so it’s a reasonable illustration. The dataset the authors of that paper provide online has the full likelihoods in terms of the flux in each energy bin, so it’s possible to construct the limits for any final state (or set of branching ratios) that you want.

Q19. Imagine there is an annihilation mechanism, but also there are DM subhalos. I set up a telescope to look at the galactic centers, e.g., but first I point it at the dark sky to calibrate the spectrum. When I build my background model (by pointing them at dark sky,) how do I avoid characterizing possible dark matter subhalos in my scope as background, which would then mess up the whole experiment?

A: This is a great question! Indeed analyses usually don’t account for this, to my knowledge. In fact sometimes analyses don’t even account for the inevitable dark matter signal from the smooth galactic halo in their “off” region, setting substructure aside - for this reason, some Galactic Center analyses lose sensitivity completely in the presence of a cored density profile, the constraint relies entirely on the assumption that the signal is rising steeply toward the Galactic Center and thus any signal in the control region can be ignored. I think the justification is that we can estimate how much substructure *should* be there from N-body simulations (albeit with a wide range), and for most searches we wouldn't expect it to be large compared to the overall astrophysical background, but this is indeed a source of systematic uncertainty.

Q20 (slide 18). Since fermi-lat observes gamma rays, so when you show constraints for bb or tau tau, does that corresponds to the photons emitted as bremstralung radiation from these?

A: No, it's mostly from the neutral pions formed when the b's hadronize and then decay, or when the tau's decay - even though the tau is a lepton, its dominant decays involve pions. (There can be contributions from internal brem, final state radiation, etc, but these are usually - but not always - suppressed, due to the extra alpha factor.) See the discussion on slide 8.

Q21 (slide 21). What causes some of these constraints to be less trusted?

A: I'm going to point you to the Carr et al review here (https://arxiv.org/abs/2002.12778), as I'd just be repeating what they say - but the dotted lines are constraints that have been suggested, but subsequent work argued that they weren't valid (in some cases the original authors have agreed with the critique, but I'm not sure this is true for all cases).

Q22 (slide 38). On the y-axis, what does a fraction of 10 mean...?

Assuming you meant page 39 in the pdf, I think they just go up to 10 so you can see clearly where the constraints cross the f=1 line - the region with f > 1 is ruled out by overclosure of the universe, so it's not physically interesting. That said, showing results up to f=10 does let you see how changes in the sensitivity to certain constraints would improve their mass reach - e.g. because the slope of the red-line constraints is so steep, even increasing their sensitivity by a factor of 10 (so their f=10 constraint would move down to be a f=1 constraint) would barely change their mass reach.

Q23 (slide 22). In this slide, in one plot you look at FSR signals from electron, and in another you directly talk about e+ e- signals, can you please explain these cases again?

Sorry, this is just misleading labeling, both are searches for the FSR photons (the e+e- plot has some other constraints on it, from other searches, which are literally looking for the electrons, but we'll talk about those tomorrow!)

Q24 (slide 23). On your PBH plot, can you clarify when you said "constrain" PBHs, do you mean that they cannot exist at those masses, or that they cannot account for the DM density in the universe? I.e. are they excluded from existing altogether, or simply excluded from being the dark matter in your decay-possible model?

f is the fraction of DM in the universe that can be comprised of PBHs, so if f=1 is excluded, then PBHs as 100% of DM are excluded, but smaller f's may still be allowed. None of these limits will constrain them from existing altogether (i.e. no limits go down to f=0).

Q25 (slide 49). How close is "neighbourhood of earth" for cosmic rays?

We'll talk about this next time - for most instruments it means the space station, at best (this is where AMS-02 is located), but Voyager is now out of the solar system and this allows it to place some surprisingly strong bounds despite being a 1970s instrument

There's a related question, which is how far do cosmic rays propagate in the galaxy (i.e. what volume are we averaging over when we measure the local cosmic-ray spectrum). Protons and antiprotons typically diffuse around the galaxy for a long time before losing all their energy, so they sample quite a large volume; high-energy electrons and positrons lose energy much faster (they are in the loss-dominated regime I mentioned), and GeV-TeV electrons and positrons typically only sample the region within about 1kpc of us.

Q26 (slide 32). What is the best way to account for solar modulation when observing proton/antiproton fluxes?

This paper from AMS-02 looks relevant:
https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.121.051101
It seems they can now measure the solar modulation in considerable detail. I have not dug deeply into this question although you're right that it's important. There is this recent paper claiming an extended force-field model can work well:
https://arxiv.org/abs/1909.01154
That said, I don't have a good handle on what the competing approaches are at present.

Q27. A general question: could you list some established experimental results of the mentioned indirect search of dark matter?

The second half of the talk is about results of various indirect searches so I'm not quite sure what you're asking about. If you mean established positive detections of interactions between dark and visible matter, then no such results yet exist (believe me, if we had confirmed non-gravitational dark matter - baryon interactions experimentally, this would be a very different summer school!) If you mean established experimental exclusions of dark matter scenarios, the slides provide many examples. There have also been discoveries in other areas of astrophysics made in the course of indirect dark matter searches - the one I'm most familiar with is the discovery of the Fermi Bubbles, https://arxiv.org/abs/1005.5480 , which I think is pretty established at this point. If you clarify your question in tomorrow's Q&A, I'll be happy to answer further.